

REMARKS

Claims 1-25 are currently active.

Claims 17-25 have been added. Antecedent support for Claims 17-25 is found in Claims 7-16, and 2-6, respectively.

The Examiner has indicated that Claims 5, 6 and 10-16 are allowable if rewritten in independent form with all the limitations of their base claim and any intervening claims. Claim 23 is Claims 2-5 written as such, and Claim 17 is Claims 7-10 written as such.

The Examiner has objected to the drawings. Substitute and corrected drawings are included to obviate this objection.

The Examiner has objected to Claim 1 for various informalities. Claim 1 has been amended consistent with the Examiner's suggestions.

The Examiner has rejected Claims 2-4 and 7 as being anticipated by Reuter1. Applicant respectfully traverses this rejection in view of the amendment to the claims.

Antecedent support for the amendments to Claims 1, 2 and 7 is found on page 6, lines 25-27.

Referring to Reuter1, there is disclosed a system and method for effectuating distributed consensus among members of a processor set in a multiprocessor computing system through the use of shared storage resources. Reuter1 teaches, with reference to figure 2, competing processors 202, 204, 206, and 208 represent processors having coordinated access to shared storage location 222. Processor 202 is associated with processor identifier 214, that uniquely identifies the processor from all other processors having access to the shared storage location 222. Likewise, processor 204 is associated with ID 216, processor 206 is associated with ID 218, processor 208 is associated with ID 220. See column 5, line 55-column 6, line 3.

Reuter1 teaches the shared storage location comprises a single unit of storage, a block of storage units, and multiple blocks of storage locally or remotely stored within the system. Reuter1 specifically teaches that the shared storage location comprises a magnetic disk drive. The shared storage location is associated with two critical storage blocks. The critical storage block is a data structure such as a register, table, or other memory or storage device. The shared storage location 222 is associated with the critical storage block 1 (CS1) 242 and the critical storage block 2 (CS2) 244. The CS1 242 preferably comprises processor

identifier field 226 and current counter field 228. The critical storage blocks provide not only means for storing state information to reserve for a mutual exclusion lock, but also a processor id field that stores the processor id. The processor id field enables a processor to detect not only the critical storage block is reserved, but identifies a specific processor that has made the reservation. See column 6, lines 16-40.


The critical storage block includes a counter field that stores a counter value to provide correctness in cases where a processor crashes and restarts the process. The counter value is incremented each time a specific processor reserves the critical storage block. The CS1 and CS2 sections may be read from and written to by any processor in a shared set of competing processors according to a mutual exclusion protocol. See column 6, lines 49-66 of Reuter1. By coordinating accesses to the CS1 242 and the CS2 244, a processor can establish an exclusive access reservation that can be detected and honored by other competing processors, without data corruption or unacceptable delays relating to race conditions or processor crashes. The programmable storage location 250 stores a race wait delay value, and storage allocation 252 stores a crash detect delay value. Each value is a time value or a clock cycle count. Each processor is associated with the same place settings as every other processor that enters the set of competing processors, although delay settings for individual processes may be tuned.

As is evident from the above, Reuter1 fails to teach or suggest the added limitation of "a set of the disk blocks are used as a communication medium", as found in amended Claim 2. As stated in the specification on page 2, lines 5-10 of the above-identified patent application, by using the disk as a communications mechanism, rather than a separate communication path, the present invention reduces the chances that the communication path used for determining which server should access the disk will fail independently from the communication path used for accessing the disk itself. Not only does Reuter1 fail to teach or suggest this data limitation, but Reuter1 fails to even recognize the problem and the advantage that this limitation added to the claims solves. Accordingly, Claim 2 is not anticipated by Reuter1. Claims 3 and 4 are dependent to parent Claim 2 and are patentable for the reasons Claim 2 is patentable over Reuter1. Claim 7 is patentable over Reuter1 for the reasons Claim 2 is patentable.

The Examiner has rejected Claims 1, 8 and 9 as being unpatentable over Reuter1 in view of Reuter2. Applicant respectfully traverses this rejection in view of the amendments to the claims. The reference Reuter2 has been cited by the Examiner simply for the reason that explains the voting algorithm taught by Reuter1 in greater detail. In pertinent part, it does not add anything in regard to teaching or suggesting the newly added limitation to Claim 1 and Claim 7. Accordingly, Claim 1 is not obvious over the applied art of record.

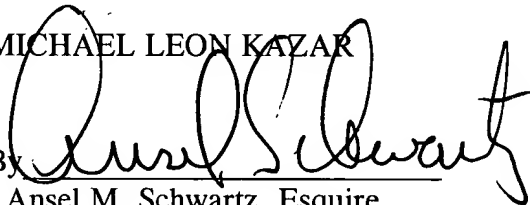
Similarly, Claim 7 is not obvious over the applied art of record. Claims 8 and 9 are dependent to parent Claim 7 and are patentable for the reasons Claim 7 is patentable.

In view of the foregoing amendments and remarks, it is respectfully requested that the outstanding rejections and objections to this application be reconsidered and withdrawn, and Claims 1-25, now in this application be allowed.



Respectfully submitted,

MICHAEL LEON KAZAR

By 

Ansel M. Schwartz, Esquire
Reg. No. 30,587
One Sterling Plaza
201 N. Craig Street
Suite 304
Pittsburgh, PA 15213
(412) 621-9222

Attorney for Applicant

ANNOTATED MARKED UP DRAWING

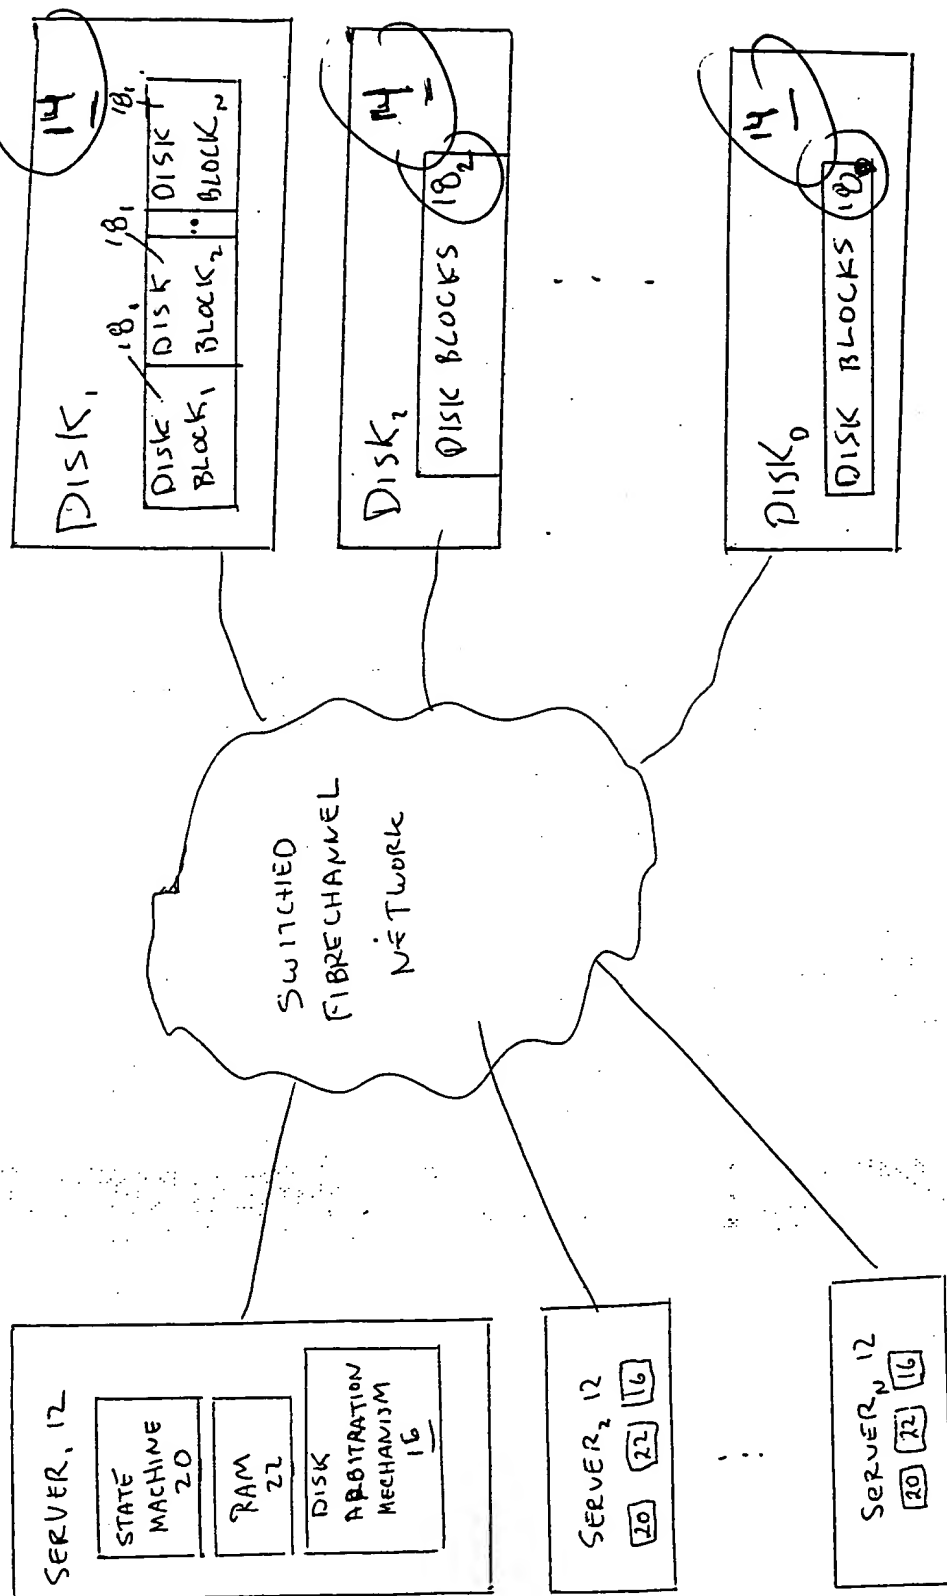


FIG 1



ANNOTATED MARKED UP DRAWING




FIG 2




FIG 3



REPLACEMENT SHEET




FIG 2




FIG 3

REPLACEMENT SHEET

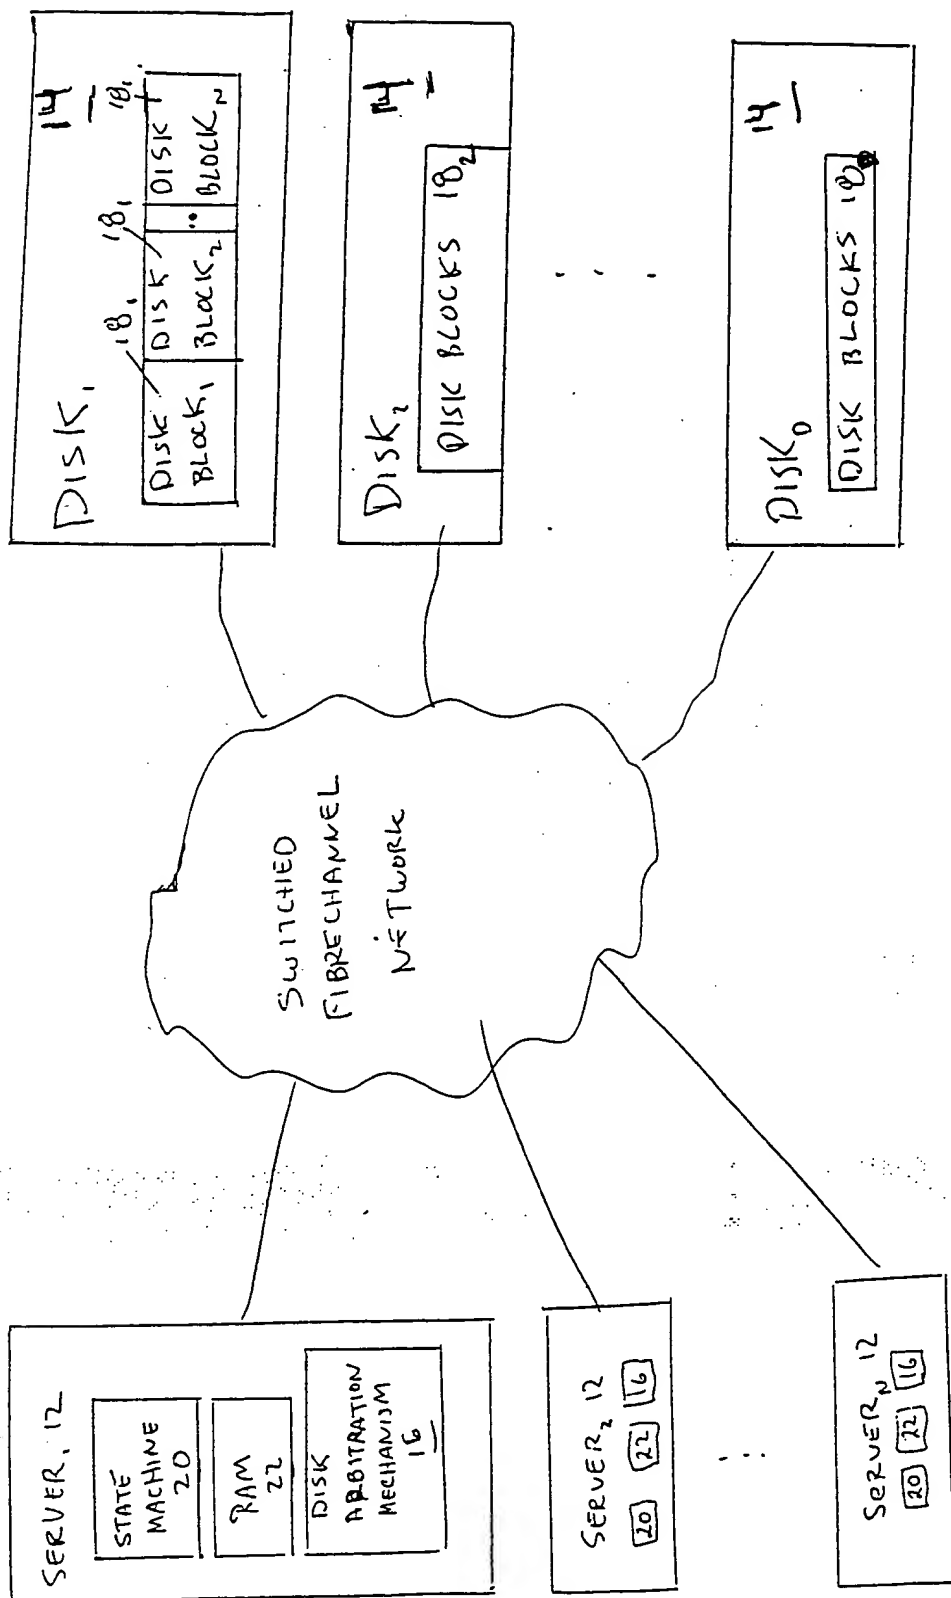


FIG 1

9

10





FIG 4

